

Income Shocks and their Transmission into Consumption*

Edmund Crawley Alexandros Theloudis

April 22, 2024

Measuring how household consumption responds to income shocks is important for understanding how families cope with adverse events, for designing government insurance or other income support policies, and for understanding the transmission of business cycles and monetary policy. It is also important for evaluating the effects of fiscal or labor market reforms on consumer welfare, and for how these reforms may impact the macroeconomy given that consumption is a large share of GDP.

This article reviews the economics literature of, primarily, the last 20 years, that studies the link between income shocks and consumption fluctuations at the household level. We identify three broad approaches through which researchers estimate the consumption response to income shocks: 1.) structural methods in which a fully or partially specified model helps identify the consumption response to income shocks from the data; 2.) natural experiments in which the consumption response of one group who receives an income shock is compared to another group who does not; 3.) elicitation surveys in which consumers are asked how they expect to react to various hypothetical events. None of these approaches is exclusive to a single field within economics; studies that use any of these methods are ordinarily classified, depending on their specific focus, in macroeconomics, labor economics, public finance – to name only a few fields.

Our aim in this short article is to survey this increasingly busy literature and provide an accessible summary of the various estimates of the consumption response to income shocks. We concentrate on the similarities and differences between the various studies, in particular with respect to the method, data, consumption notion, and type of income shock analyzed, thus also with respect to the type of consumption response each work identifies. Our focus

*Crawley: Federal Reserve Board; email: edmund.s.crawley@frb.gov. Theloudis: Department of Econometrics & OR, Tilburg University; email: a.theloudis@gmail.com. Viewpoints and conclusions stated in this article are the responsibility of the authors alone and do not necessarily reflect the viewpoints of the Federal Reserve Board. This review will appear in the Encyclopedia of Consumption.

is on responses to shocks, i.e., *unanticipated* income changes; [Jappelli and Pistaferri \(2010\)](#) review the earlier evidence on responses to *anticipated* income changes.

The survey proceeds as follows. Section 1 introduces a brief theoretical framework that helps fix ideas for the subsequent discussion. The next sections are devoted to the different approaches to the estimation of the consumption response. Section 2 surveys the studies that employ structural methods, section 3 reviews the evidence from natural experiments, and section 4 focuses on the elicitation surveys. Section 5 concludes.

1 Theoretical Background

A household i chooses consumption C_{it} , savings A_{it+1} , and perhaps other behaviors captured in L_{it} (possibly a vector), e.g., labor supply, to maximize its expected lifetime utility

$$\max_{\{C_{it}, A_{it+1}, (L_{it})\}_{t=0}^T} \mathbb{E}_0 \sum_{t=0}^T \beta^t U_i(C_{it}; L_{it}), \quad (1)$$

for which we assume separability over time and geometric discounting.¹ Expectations about future states of the world are captured by \mathbb{E}_0 . The utility function depends on consumption and on L_{it} , which may be a choice variable (e.g., endogenous labor supply) or taken as given (e.g., exogenous labor supply). Utility is subscripted by i to reflect preference heterogeneity across households. We have employed a finite-horizon setting with T as its terminal period, which is natural given our focus on households.²

The problem is subject to the sequential budget constraint

$$(1+r)A_{it} + T_i(Y_{it}^g; L_{it}) = C_{it} + A_{it+1}, \quad (2)$$

which links resources over time under the assumption that consumers can borrow and save at an interest rate r .³ T_i maps before taxes/transfers gross income Y_{it}^g to disposable household income Y_{it}^d . Y_{it}^g may be a vector; for instance, with two earners in the household, $Y_{it}^g = (Y_{it}^{e1}, Y_{it}^{e2})'$, where Y^{e_j} indicates earnings of member j . T_i may depend on choices L_{it} ; for example, with endogenous labor supply, the primitive source of gross income is the hourly wage, i.e., $Y_{it}^g = (Y_{it}^{w1}, Y_{it}^{w2})'$, where Y^{w_j} indicates the wage of member j . T_i is subscripted by i to reflect heterogeneity in taxes, welfare benefits, contingent transfers, or external sources of income (what is often called external insurance). There may be a borrowing constraint in

¹Strictly, the household chooses a plan for C_{it} , A_{it+1} , and L_{it} contingent on future states of the world, including, e.g., future states of income. We do not explicitly show the contingent states to ease the notation.

²The extension to infinite horizon is trivial and mostly inconsequential ([Jappelli and Pistaferri, 2010](#)).

³The extension to multiple or risky assets is also straightforward.

some periods such that $A_{it} \geq B_{it}$, where B_{it} denotes the applicable borrowing limit. Finally, there is a terminal condition on A_{iT+1} , which reflects that households run down their assets before death or bequeath their offspring.

There are many alternative specifications for the process that governs income; a general formulation is

$$Y_{it}^k = f_i^k(X_{it}, Y_{it-1}^k, v_{it}^k, u_{it}^k, \dots), \quad (3)$$

where $k = \{d, g, e_1, e_2, w_1, w_2, \dots\}$ indicates the type of income considered, i.e., disposable, gross, etc. f_i^k reflects the precise process, which allows for heterogeneity across households and depends on observables X_{it} (e.g., age, time, education), past income Y_{it-1}^k , and idiosyncratic shocks v_{it}^k and u_{it}^k to *log* income, such as permanent and transitory shocks. It may also depend on older income, other shocks including aggregate ones, and past shocks, depending on the specific process that is being considered. [Meghir and Pistaferri \(2011\)](#) offer a review of the vast income dynamics literature and discuss popular cases.⁴

Solving (1) subject to (2), the borrowing constraint, and the terminal condition, yields a consumption policy rule whose exact formulation depends on the preference specification, the income process (3), the tightness of the borrowing constraint, and the market environment in which the household operates (for example, the extent to which it has access to contingent transfers or external insurance). A general formulation for the consumption rule is

$$C_{it} = g_i(A_{it}, X_{it}, Y_{it-1}^k, v_{it}^k, u_{it}^k, \dots), \quad (4)$$

which allows for heterogeneity across households (reflecting, among other things, heterogeneity in U_i) and generally depends on assets and the various components of income. Policy rule (4) subsumes several popular settings in the literature, and [Jappelli and Pistaferri \(2010\)](#) offer specific examples. It is nonetheless non exhaustive of all possible settings. For instance, durable goods necessitate accounting for their stock and possible adjustment costs. Here we simply see (4) as a general organizational device rather than the solution to any given model.

Interest lies in how and by how much the income shocks v_{it}^k , u_{it}^k , etc, impact consumption. To measure this, the literature focuses on two main parameters, the marginal propensity to consume (MPC), broadly defined as the derivative of consumption with respect to income – i.e., dC_{it}/dY_{it}^k , and the pass-through rate, broadly defined as the derivative of consumption growth with respect to the shock – i.e., $d\Delta \ln C_{it}/dv_{it}^k$. We now turn to the three broad approaches to estimating these parameters.

⁴There are multiple generalizations of (3). For example, one may allow f_i^k to depend on time (age), thus enable the impact of shocks to be time-varying. We view (3) as a simple organizing device rather than as an exhaustive representation of every possible income process.

2 Structural Methods

Early papers. A series of influential papers in the 1980s-90s test the predictions of the permanent income hypothesis and the complete-markets model, the then benchmark models in the literature. This is done through forming appropriate hypotheses on the link between consumption and income fluctuations that emanate from these models. While this early work does not strictly *measure* MPCs or pass-through rates, it provides motivating evidence for the subsequent work that explicitly measures the consumption response to shocks.

Hall and Mishkin (1982), one of the first studies, investigate the sensitivity of food consumption to income using micro data from the Panel Study of Income Dynamics (PSID). Consumption varies more closely with permanent than with transitory shocks, which is an implication of the permanent income hypothesis in which households smooth consumption through self-insurance (saving, borrowing). Yet, the sensitivity of consumption to transitory income is much stronger than theory predicts, which leads to rejection of the model.⁵

The permanent income hypothesis postulates that inequality in consumption grows over the lifecycle. This is the motivating observation for Deaton and Paxson (1994) who, using micro data from multiple countries, confirm that the variance of consumption (and income) grows with age. While they cannot reject the permanent income hypothesis, they admit that the evidence is also consistent with other models of intertemporal choice, such as models that permit some external insurance to idiosyncratic income shocks.

On the opposite end of theory, Cochrane (1991) tests for full insurance by assessing the sensitivity of consumption to income growth and to events such as illness and job loss. Under complete markets, consumers have access to contingent transfers, so household consumption growth should be unrelated to *idiosyncratic* events. Focusing on food consumption in the PSID, the hypothesis is rejected following long illness or job loss, but not rejected in response to short unemployment spells, thus providing early evidence for partial insurance.⁶

Attanasio and Davis (1996) test for full insurance across birth cohorts and education groups, accounting for consumption-work complementarity and common demographics driving consumption and wages. Under complete markets and absent of aggregate shocks, consumption growth should not co-move with wage growth. Drawing synthetic panels from the Consumer Expenditure (CEX) and Current Population (CPS) surveys, they sharply reject full insurance to low-frequency wage shifts. In a similar model, Hayashi et al. (1996) test for full insurance across and within extended families, which they reject based on the strong

⁵Other early papers that test the permanent income hypothesis are Hall (1978), who does not reject it, Sargent (1978), who rejects it, and Flavin (1981), who also rejects it. They all use time series data.

⁶Altug and Miller (1990) model a complete-markets environment and allow for non-separability with labor supply; using food consumption in the PSID, they cannot reject full insurance to wage fluctuations.

correlation between wage and food consumption growth in the PSID. Using similar data in a simpler setting, [Altonji et al. \(1992\)](#) also reject full insurance. By contrast, [Mace \(1991\)](#) and, in particular, [Townsend \(1994\)](#), who studies village insurance in India, find mixed evidence.

In sum, these early works frequently reject the benchmark models of permanent income (self insurance) and complete-markets (full insurance). Yet, the data consistently reveal that households have access to *some* insurance to income shocks.⁷ The literature in the busy 2000s-10s attempts to measure the degree of insurance and identify its sources.

Covariance restrictions. [Blundell et al. \(2008\)](#), abbreviated as BPP, introduce the seminal methodology to measure the consumption response to income shocks. Their idea is that the extent to which consumption growth varies with income growth, the latter being driven by various income shocks, reflects the *degree* of transmission of those shocks into consumption.⁸ This is motivated by a consumption process that is log-linear in income shocks, namely

$$\Delta c_{it} = \xi_{it} + \phi_t v_{it}^k + \psi_t u_{it}^k, \quad (5)$$

where Δc_{it} is consumption growth $\Delta \ln C_{it}$ net of observables X_{it} .⁹ v_{it}^k and u_{it}^k are, respectively, a permanent and a transitory shock to disposable household income (so $k = d$ here), ϕ_t and ψ_t are their transmission parameters, and ξ_{it} is a preference shock unrelated to income. If income follows the canonical permanent-transitory process, namely $\Delta y_{it}^k = v_{it}^k + \Delta u_{it}^k$ where Δy_{it}^k is income growth $\Delta \ln Y_{it}^k$ net of X_{it} , the transmission parameters are identified through $\text{Cov}(\Delta c_{it}, Z_{it}) / \text{Cov}(\Delta y_{it}^k, Z_{it})$, namely a regression of Δc_{it} on Δy_{it}^k using appropriate instruments Z for income. In the case of permanent shocks, permanent income $Z_{it} = \sum_{\tau=-1}^1 \Delta y_{it+\tau}^k$ nets Δy_{it}^k from the transitory shock at t , so its covariance with Δc_{it} identifies ϕ_t . In the case of transitory shocks, future income $Z_{it} = \Delta y_{it+1}^k$ shifts Δy_{it}^k due to mean reversion of the transitory shock, so its covariance with Δc_{it} identifies ψ_t . It is precisely these covariance restrictions that this strand of literature has taken its name from.

The consumption equation (5) can be obtained through a log-linearization of the policy rule (4) in a lifecycle permanent income model with CRRA utility, a permanent-transitory income process, and slack borrowing constraints. In this case, ϕ_t reflects the share of consumption that is funded by future labor income (as opposed to assets whose value remains

⁷Among the first studies to quantify the extent of consumption insurance, [Gruber \(1997\)](#) measures how UI reduces the fall in consumption upon unemployment.

⁸In a predecessor paper, [Blundell and Preston \(1998\)](#) assume that that the rise in consumption inequality observed by [Deaton and Paxson \(1994\)](#) is driven by permanent but not transitory shocks, to which consumers can fully self-insure. This allows them to use income and consumption moments to identify the variances of permanent and transitory income shocks. [Attanasio et al. \(2002\)](#) extend this to a setting of two earners with separate income streams, while [Primiceri and van Rens \(2009\)](#) extend it to heterogeneous income processes.

⁹We use Δ to denote the first difference operator; so $\Delta X_t = X_t - X_{t-1}$.

unchanged by the shock to income) and, depending on the measure of Y_{it}^k , features of the tax/benefits system; ψ_t is similar up to an annuitization factor for the household’s remaining horizon. Yet, (5) is also consistent with other settings in which the transmission of shocks depends on their persistence, e.g., environments with moral hazard (Attanasio and Pavoni, 2011) or external insurance (Blundell et al., 2016). The linear insurance equation (5) is, therefore, the *reduced form* of multiple environments that may differ in their insurance content. As such, ϕ_t and ψ_t are called *partial insurance* parameters; they measure the *overall* pass-through of (or, on the flip side, insurance to) permanent and transitory income shocks, regardless of the precise mechanisms that give grounds for such insurance.

Empirically, BPP focus on nondurable consumption and disposable household income in 1980-92. While the PSID is ideal for its income data, its consumption data, including mostly food items, is limited until 1999; BPP thus impute consumption from the CEX into the PSID. They estimate $\phi_t = 0.64$ (a 10% permanent income cut reduces consumption by only 6.4% – households are thus partially insured to permanent shocks) and $\psi_t = 0.05$ (statistically not different from zero, so consumption is fully insured against transitory shocks).

BPP sparked multiple extensions. Blundell et al. (2013) derive (5) under an autoregressive income process and general preferences. They assess the approximation vis-à-vis the true policy rule and show that it performs well when liquidity constraints do not bind. Blundell et al. (2016) model endogenous labor supply to measure the consumption response to *wage* shocks. Using income and consumption data in the PSID after 1999, they estimate $\phi_t = 0.32/0.19$ (male/female wages) and find little external insurance when family labor supply, assets, and the tax/benefits system are accounted for.¹⁰ Theloudis (2017) extends this to an intra-household bargaining (collective) setting, in which lack of commitment to lifetime marriage limits the insurance role of family labor supply. Chopra (2023) argues that the insurance role of family labor supply increases during recessions.

Most of these works permit limited heterogeneity. Arellano et al. (2017), ABB in short, relax linearity in the income process and let shocks feature *nonlinear* persistence depending on their sign and size. The consumption response to shocks varies flexibly with their level and past history. Using a quantile-based estimation method and PSID data over 1999-2011, they measure the consumption response to log income, a type of MPC, at 0.2-0.4 on average (specification with household heterogeneity), though the response varies over the distribution of shocks. Although allowing for heterogeneity makes interpreting their parameter similar to a pass-through rate, the response is markedly different from BPP. This may be due to

¹⁰Hyslop (2001) uses covariance restrictions to measure the pass-through of wage shocks into household earnings. Jessen and König (2023) is related: using hours and earnings data in the PSID over 1970-1997, they find that wage and taste shocks have a comparable contribution to the total variance of earnings.

ABB’s flexible income process or the new consumption data in the PSID.¹¹

Ghosh and Theloudis (2023) relax linearity in the consumption process by writing Δc_{it} as a quadratic polynomial in income shocks. This stems from a second-order approximation to the policy rule (4) in a model similar to BPP. The pass-through of shocks now depends on their sign and size, so this method allows small vs. large, or good vs. bad shocks, to have an *asymmetric* impact on consumption. Identification requires second, third, and fourth moments of income and second moments of consumption, in contrast to ABB who require observing their entire distribution. Using PSID data after 1999, they estimate the pass-through of the *average* permanent shock at 0.13; bad or large permanent shocks have a much bigger impact on consumption, and their pass-through increases with their severity.¹²

Alan et al. (2018) estimate the extent of insurance to income shocks allowing for flexible joint heterogeneity in the consumption and income processes. They use PSID data over 1968-2009 and find pervasive cross-household heterogeneity in the pass-through, ranging from 0.05 to 0.69. Theloudis (2021) allows for unobserved preference heterogeneity in a model with family labor supply. He explores higher-order moments of income and consumption to identify the contribution of heterogeneity to consumption inequality.

A consistent empirical finding in this literature is that, on average, consumption is fully insured against transitory shocks. Several papers challenge this result. Crawley (2020) argues that BPP neglect time aggregation in the PSID. Time aggregation occurs when income is observed less frequently (annually) than the underlying true data (e.g., monthly payments); income growth is then *mechanically* positively correlated, which changes the covariance restrictions used to identify the variance of shocks and their transmission into consumption. Crawley (2020) addresses this point and estimates the pass-through of shocks at $\phi_t = 0.34$ and $\psi_t = 0.24$ (statistically significant).

Commault (2022) extends the linear insurance model (5) to allow consumption to respond to *past* transitory shocks. There are several underlying structures that justify this. The exact solution to Δc_{it} in a lifecycle permanent income model includes higher-order terms that depend on past variables (e.g., wealth); models with limited commitment is another example. Commault (2022) proposes a new estimator for ψ_t , one that selects as instrument the only future $\Delta y_{it+\kappa+1}^k$ that is uncorrelated with past shocks. She estimates $\psi_t = 0.6$ (MPC at 0.32), which helps bridge the discrepancy in ψ_t between studies that employ covariance restrictions and those that rely on natural experiments.¹³

¹¹ABB is not a covariance method, but their framework falls firmly in this category because they measure the *overall* impact of shocks, as in BPP. Arellano et al. (2024) advance this to unbalanced panels and flexible heterogeneity. They measure the consumption response to log income at 0.2 on average.

¹²The average pass-through is close to ABB but markedly different from BPP. The discrepancy is due to the consumption imputation in BPP and the biennial (vs. annual) frequency of the modern data.

¹³As we review subsequently, natural experiments often imply a larger ψ_t than typically estimated through

Crawley and Kuchler (2023) address neglected time aggregation *and* allow consumption to depend on past transitory shocks. They estimate heterogeneous ϕ_t and ψ_t over well-defined subpopulations using Danish administrative data, and relate their estimates to liquid wealth and other household balance-sheet characteristics.

Hryshko and Manovskii (2022) identify households in the PSID with vastly different degrees of insurance. They show that the *sons* of families originally surveyed by the PSID in 1968 exhibit almost no insurance, while the *daughters* have substantial partial insurance. They do a thorough job explaining this discrepancy by differential income persistence across the two groups, which is further explained by differential attrition from the survey.

Fully specified models. The studies employing covariance restrictions do not *fully* specify preferences, expectations, and the budget set, so they take no stance on the exact mechanisms that give rise to partial insurance. Due to their semi-parametric nature, these works are of limited use for policy counterfactuals. Another approach is to fully specify the channels through which consumers smooth income shocks, and take such models to the data.

Attanasio et al. (2005) quantifies the insurance role of female labor supply (measured in terms of welfare costs of income uncertainty) through a structural model of consumption and labor supply with earnings risk. Krueger and Perri (2006) show that a model with a complete set of Arrow securities but limited enforceability of contracts reproduces income and consumption inequality in the U.S., suggesting that households possess more insurance than self-insurance. Storesletten et al. (2004) show that a lifecycle permanent income model can also produce empirically consistent income and consumption inequality if the tax/benefits system and the aggregate level of wealth are accounted for. Heathcote et al. (2008) calibrate a lifecycle model of consumption/labor supply with partial insurance to wage shocks to measure the welfare costs of risk and market incompleteness. Partial insurance is fixed by the authors: permanent shocks are uninsured, while transitory shocks are fully insured. Low et al. (2010) calibrate a lifecycle model of consumption, labor supply, and job mobility with income and employment risk. Self-insurance aside, the model allows for three channels of partial insurance: unemployment benefits, disability insurance, and food stamps.

These earlier quantitative models measure the welfare implications of risk (or of certain insurance mechanisms) but not the degree of consumption insurance per se.¹⁴ Kaplan and Violante (2010) explicitly measure the degree of insurance in a calibrated lifecycle per-

covariance restrictions. It is unclear, however, if survey data such as the PSID, on which the latter studies rely, reflect larger transitory shocks that are typical of experiments. Studies that allow the pass-through to depend on the size of the shock typically find that larger shocks have larger pass-through, thus also helping bridge the gap in estimates between natural experiments and covariance restrictions.

¹⁴An exception is Low et al. (2010) who report the consumption response to an unemployment shock.

Table 1: Summary of Estimates from Structural Methods

| Study | Pass-through | | MPC | | Variables | | Data |
|---|--------------|---------|---------|----------|------------|----------|---------------------------------|
| | perm. | trans. | perm. | trans. | Y_{it}^k | C_{it} | |
| Alan et al. (2018) | 0.05–0.69 | | | | thy | food | P 1999-2009 |
| Arellano et al. (2017) ^a | | | 0.2–0.4 | -0.4–0.2 | thy | nde | P 1999-2009 |
| Arellano et al. (2024) ^b | | | 0.33 | | dhy | nde | P 2005-17 |
| Blundell et al. (2008) | 0.64 | 0.05 | | | dhy | nde | P & C 1980-92 |
| Blundell et al. (2016) | 0.32 | -0.14 | | | mhw | nde | P 1999-2009 |
| | 0.19 | -0.04 | | | fhw | nde | P 1999-2009 |
| Blundell et al. (2018) | 0.39 | 0.12 | | | mhw | nde | P 1999-2015, and |
| | 0.35 | 0.13 | | | fhw | nde | C & A 2003-15 |
| Busch and Ludwig (2021) ^c | 0.40 | 0.05 | 0.38 | 0.05 | dhy | n/a | P 1977-2012 |
| Chopra (2023) | 0.29 r | -0.18 r | 0.19 r | | mhw | nde | P 1977-2016 |
| | 0.31 x | -0.26 x | 0.12 x | | mhw | nde | P 1977-2016 |
| Commault (2022) | | 0.6 | | 0.32 | dhy | nde | P & C 1980-92 |
| Crawley (2020) | 0.34 | 0.24 | | | dhy | nde | P & C 1980-92 |
| Crawley and Kuchler (2023) | | | 0.64 | 0.64 | dhy | te | D 2003-15 |
| De Nardi et al. (2019) | 0.54 | 0.12 | | | dhy | nde | P 1968-92, and C 1980-2007 |
| Ghosh and Theloudis (2023) ^c | 0.13 | -0.00 | | | dhy | nde | P 1999-2019 |
| Güvener and Smith (2014) | | 0.45 | | | dhy | nde | P & C 1968-92 |
| Güvener et al. (2023) ^d | 0.38 | 0.11 | 0.4 | 0.05 | dhy | nde | external estims |
| Heathcote et al. (2014) | 0.39 | | | | mhw | nde | P 1968-2007, and C 1980-2006 |
| Hryshko and Manovskii (2022) | 0.87 sn | 0.07 sn | | | dhy | nde | P & C 1980-92 |
| | 0.46 dg | 0.12 dg | | | dhy | nde | P & C 1980-92 |
| Jessen and König (2023) | 0.62 | | | | mhw | n/a | P 1970-1997 |
| Kaplan and Violante (2010) | 0.78 | 0.06 | | | dhy | nde | P 1980-92, SCF |
| Low et al. (2010) ^e | 0.56 | | | | mhw | n/a | P 1988-96, SIPP |
| Madera (2019) ^d | 0.50 | 0.10 | | | dhy | te | P 1999-2015 |
| Theloudis (2021) | 0.45 | -0.03 | | | mhw | nde | P 1999-2011 |
| | 0.27 | -0.05 | | | fhw | nde | P 1999-2011 |
| Wu and Krueger (2021) | 0.35 | 0.01 | | | mhw | nde | P 1999-2009 |
| | 0.18 | 0.01 | | | fhw | nde | P 1999-2009 |

Legend: **dg**: daughters; **dhy**: disposable household income; **fhw**: female hourly wage; **mhw**: male hourly wage; **nde**: non-durable expenditure; **n/a**: not applicable; **r**: recession; **sn**: sons; **te**: total expenditure; **thy**: total household income; **x**: expansion; **A**: American Time Use Survey; **C**: Consumer Expenditure Survey; **D**: Danish registry data; **P**: Panel Study of Income Dynamics; **SCF**: Survey of Consumer Finances; **SIPP**: Survey of Income & Program Participation.

^aResults with unobserved household heterogeneity, figures S21 and S24.

^bResults with filtering and unobserved household heterogeneity.

^cResults for average/medium shock.

^dResults at age 40.

^ePass-through of income following an unemployment shock.

manent income model with income risk. This is the model whose consumption rule BPP log-linearize. While they find almost full insurance to transitory shocks, as in BPP, they estimate $\phi_t = 0.78$, larger than BPP's estimate. Households in the model possess *less* insurance to permanent shocks than in the data, highlighting that *self-insurance*, the only mechanism

in the model, is not enough to generate the excess consumption smoothing we observe empirically. By contrast, [Wu and Krueger \(2021\)](#) calibrate a lifecycle permanent income model with endogenous labor supply and find that it matches the empirical pass-through rates of male and female *wage* shocks in [Blundell et al. \(2016\)](#). Family, and mainly *female*, labor supply is a crucial insurance mechanism that previous studies had neglected.

[Guvenen and Smith \(2014\)](#) estimate a lifecycle consumption–savings model with self-insurance, external insurance, and learning over stochastic income. They use the joint dynamics of earnings and consumption in the PSID and CEX to quantify earnings risk and the extent of insurance to it. About half of the earnings shock (including permanent and transitory elements) is smoothed. [Heathcote et al. \(2014\)](#) estimate a general equilibrium consumption/labor supply model with partial insurance to wage shocks. They model self-insurance, labor supply, a tax/benefits system, and external insurance. In their benchmark, they use PSID earnings and hours data alone and find that 39% of permanent wage shocks pass through into consumption. [Blundell et al. \(2018\)](#) present a structural lifecycle consumption–savings model with labor supply and childcare; in this setting, childcare responds endogenously to wage shocks and acts as an additional insurance mechanism.

The previous studies assume income risk is Gaussian. [Guvenen et al. \(2021\)](#) and other authors establish that the distribution of income shocks exhibits substantial left skewness and excess kurtosis. This implies that there are far more people in the data who experience small unimportant or extreme negative shocks than people who experience moderate or extreme positive ones. A newer literature attempts to measure the pass-through of income shocks accounting for these higher-order features of income dynamics.

[De Nardi et al. \(2019\)](#) estimate a lifecycle model of consumption/savings with non-Gaussian income risk. They show that tail income risk increases the degree of partial insurance to permanent shocks due to stronger precautionary motives. [Busch and Ludwig \(2021\)](#) estimate a similar model explicitly targeting income skewness and kurtosis. They distinguish between good and bad shocks and find that the latter are worse insured than the former. In a related model, [Madera \(2019\)](#) studies the differential response of durable and nondurable consumption to tail earnings shocks. Durable consumption responds more strongly to tail shocks than nondurable consumption does. In a calibrated lifecycle consumption-savings model, [Guvenen et al. \(2023\)](#) establish that non-Gaussian earnings risk implies large welfare losses and commands a strong consumption response. They confirm that the benchmark method in BPP understates the true consumption response to such shocks, as derived analytically by [Ghosh and Theloudis \(2023\)](#) in the presence of tail income risk.

3 Natural Experiments

At the other end of the spectrum to structural models are reduced-form studies of natural experiments. These studies usually compare one group of individuals or households who have received a shock to their income, such as a stimulus check, with a group that has not. In contrast to studies of structural models, natural experiment studies tend to focus on estimating a marginal propensity to consume (MPC) as opposed to a pass through parameter. This is partly because the identified shock is not often proportional to income and indeed the researchers rarely know what household income is.

There are two main challenges for a researcher estimating MPCs using natural experiments. The first is to find a suitable natural experiment and the second is to find high-quality data on spending or consumption at the individual or household level. The literature has boomed recently as more and more data sources have become available to researchers. Some more recent papers are able to speak to MPC heterogeneity across household demographics and asset holdings, as well as shock size and sign.

Overall, the evidence from these natural experiments points to larger consumption responses to income shocks than from the structural modeling literature. However, the range of estimates is wide and there is still no consensus in the profession. [Havranek and Sokolova \(2020\)](#) examine 144 studies of excess sensitivity and find evidence of publication bias that suggests MPCs may be smaller than implied by a survey of this literature such as this one.

[Johnson et al. \(2006\)](#) was one of the first papers to study a convincing natural experiment. The authors look at household spending following the 2001 tax rebates in the United States. The distribution of these rebates, typically \$300 or \$600 in size, were staggered according to the second-to-last digit of recipients' social security number—an effectively random assignment. Using the CEX, the authors find that households spent between 20 and 40 percent of their rebates on non-durable goods during the three-month period in which they received their rebates, with some evidence that households with low liquid wealth or low income had larger MPCs. [Parker et al. \(2013\)](#) use a similar research design on the 2008 stimulus checks in the United States and found a slightly smaller consumption response for nondurable goods, but a large total response—including durables—of 50 to 90 percent in the three-month period in which the check arrived.

These estimates of the spending response to U.S. stimulus policies are highly cited and often contested. For example, [Misra and Surico \(2014\)](#) use quantile regression techniques on the same two experiments and find MPC estimates that tend to be smaller and more accurate than those from a homogeneous model. [Orchard et al. \(2023\)](#), while pointing to the difficulty of reconciling high MPC estimates with the macro evidence, also find lower MPC

estimates when updating the results from [Parker et al. \(2013\)](#) with new insights into the difference-in-differences methodology used. Pointing in the other direction is evidence that households significantly under report their spending on goods and non-housing services in the CEX—see [Sabelhaus et al. \(2014\)](#). Overall, these two episodes of government fiscal stimulus have been the subject of much research, partly because the random assignment provides an excellent identification scheme and partly because the effectiveness of such stimulus policies is a question of vital importance in its own right.

[Shapiro and Slemrod \(2009\)](#) and [Sahm et al. \(2010\)](#) study the 2008 tax rebates using questions on the Michigan Survey of Consumers. They find that around 20% of respondents say that they will “mostly spend” the rebate when asked whether they would use the rebate to mostly spend, save, or pay off debt. The authors combine this with other evidence to suggest an implied MPC of around one third. Interestingly, [Sahm et al. \(2012\)](#)—who also use the Michigan Survey of Consumers—look at differences in consumption responses between a one-time payment relative to flow of payments from reduced tax withholding. They find that the one-time payment may be about twice as effective at inducing spending, highlighting some of the difficulty in pinning down consumer behavior.

During the pandemic, three Economic Impact Payments (EIPs) were distributed to households and their effects have been widely studied; [Falconetti and Nygaard \(2021\)](#) and [Gelman and Stephens \(2022\)](#) provide more thorough reviews of the related literature. These payments were significantly larger than previous economic stimulus payments, summing to a maximum of \$11,400 for a family of four over a period of less than one year. The EIPs were also distributed more quickly than similar programs in the past; as a result, they were not staggered by social security number and thus identification of their effects is not as clean as was possible in 2001 and 2008. Furthermore, they were distributed at a time of highly unusual consumption behavior, both at the aggregate level and between income groups and geographies, all of which make it harder to tease out their spending effects. Nevertheless, the EIPs have proved a fruitful source of knowledge about the effectiveness of stimulus programs.

[Parker et al. \(2022\)](#) is closely aligned with the methodology of [Johnson et al. \(2006\)](#) and [Parker et al. \(2013\)](#), making use of the CEX, but, because the randomly staggered research design is not available, this newer analysis “leans on comparing the spending of similar households that do and do not receive EIP and that receive EIPs of different amounts relative to their typical spending amounts”. They find that households spent their EIPs more slowly on average than the stimulus payments in 2001 and 2008.

Relative to the 2001 and 2008 stimulus payment studies, by 2020 many researchers had access to transaction level data. These data mitigate concerns about survey respondents under reporting their spending but can suffer more from sample selection bias than surveys

designed by statistical agencies, such as the CEX. [Karger and Rajan \(2020\)](#), [Misra et al. \(2022\)](#), and [Baker et al. \(2023\)](#) all make use of data from personal finance apps. These convenience samples skew somewhat towards low-income households, while the SaverLife app studied in [Baker et al. \(2023\)](#) is specifically aimed at helping users to save money. All these studies find relatively high MPCs within the first few weeks after the first EIP payment is received—somewhere between 0.25 and 0.5 on average. Furthermore, and in contrast to some other studies, [Baker et al. \(2023\)](#) finds little response of durables spending.

Recently, there has been a growing interest in heterogeneous agent models in macroeconomics that has increased the demand for reliable estimates of MPC heterogeneity, with a particular focus of the role of liquidity in determining MPCs. Some of the best evidence on MPC heterogeneity has come from outside of the United States. In [Fagereng et al. \(2021\)](#), the authors study the spending response of Norwegian lottery winners using administrative data. They find that, in line with buffer-stock models, MPCs are negatively correlated with the winner’s stock of liquid assets and negatively correlated with the size of the lottery win. They also find that MPCs decline with age, even controlling for liquid assets. Despite these first two correlations going in the expected directions, the high level of spending in the year in which the lottery is won, especially for households with significant liquid asset holdings, is difficult to reconcile with standard models. A further benefit of this study of Norwegian lottery winners is that the consumption response can be tracked over several years and is thus informative not just of the initial MPC but can also be used to inform the so-called intertemporal MPC—how households react to a shock to their income over time. The evidence suggests that spending is more front-loaded than a standard model would predict, especially given the large size of the lottery wins studied.

In a study of French households, [Boehm et al. \(2022\)](#) conduct a randomized controlled trial to analyze the consumption response to unanticipated one-time money transfers of €300. The trial randomized three different types of transfer via a prepaid debit card: one transfer with no restrictions, one transfer in which unspent value expired within three weeks, and one transfer subject to a 10% negative interest rate every week. The authors were able to see participants main bank accounts and therefore the effect of the transfers on their overall spending. They find an overall MPC of 0.23 within the first month of the transfer. Spending is concentrated in the first few weeks following the transfer, after which there is little boost to spending. They observe significant MPC heterogeneity by observable household characteristics including by liquid wealth, current and permanent income, and gender. However, similar to the Norwegian lottery study, they find that even households with high liquid asset holdings have high MPCs and—something that is not observable in the Norwegian data—that the spending response is concentrated in the short run for non-

Table 2: Summary of Estimates from Natural Experiments and Elicitation Surveys

| Natural experiment study | MPC | Horizon | C_{it} | Data |
|--|-------------------------|--------------|----------|-------------------------|
| Baker et al. (2023) | 0.25–0.40 | First weeks | te | SaverLife 2020 |
| Boehm et al. (2022) | 0.23 | One Month | te | French RCT 2022 |
| Fagereng et al. (2021) | 0.35–0.71 | First year | te | N 1993-2015 |
| Gelman et al. (2023) | ≈ 1.00 | Perm. shock | te | FA 2013-16 |
| Johnson et al. (2006) | 0.2–0.40 | Three months | nde | C 2001 |
| Misra and Surico (2014) | 0.43 | Three months | nde | C 2001 |
| | 0.16 | Three months | te | C 2008 |
| Karger and Rajan (2020) | 0.46 | Two weeks | te | Facteus 2020 |
| Misra et al. (2022) | 0.29 | A few days | te | Facteus 2020 |
| Orchard et al. (2023) | ≈ 0.3 | Three months | te | C 2008 |
| Parker et al. (2022) | 0.05–0.16 | Three months | nde | C 2020-21 |
| Parker et al. (2013) | 0.50–0.90 | Three months | te | C 2008 |
| Sahm et al. (2010) | ≈ 0.3 | One year | te | C 2008 |
| Elicitation survey study | MPC | Horizon | C_{it} | Data |
| Bunn et al. (2018) | 0.14 (pos.) 0.64 (neg.) | One year | te | BoE survey 2011-14 |
| Christelis et al. (2019) | 0.20 (pos.) 0.24 (neg.) | One year | te | Dutch survey 2015 |
| Colarieti et al. (2024) | 0.16 | One quarter | te | Authors’ survey 2022-23 |
| Fuster et al. (2020) | 0.07 (pos.) 0.32 (neg.) | Three months | te | NY Fed SCE 2016-17 |
| Jappelli and Pistaferri (2014) | 0.48 | Unspecified | te | SHIW 2010 |
| Jappelli and Pistaferri (2020) | 0.47 | Unspecified | te | SHIW 2016 |

Legend: **neg.:** negative; **nde:** non-durable expenditure; **pos.:** positive; **te:** total expenditure; **BoE:** Bank of England; **C:** Consumer Expenditure Survey; **FA** Financial Aggregator; **N** Norwegian registry data; **SCE:** Survey of Consumer Expectations; **SHIW:** Italian Survey of Household Income and Wealth.

durables. Finally, MPCs are highest for the group given the card that becomes unusable after three weeks and lowest for the group given the card that has no restrictions on its use. Many other studies have found correlations between MPCs and observables, particularly a negative relation between MPCs and liquid wealth. However, observables are likely to only explain a small fraction of all MPC heterogeneity, as documented by [Lewis et al. \(2019\)](#).

A standard income shock like those above mixes both the effects of an increase in lifetime budget constraint with the effects of an increase in liquidity. [Hamilton et al. \(2023\)](#) analyze a policy implemented in Australia during the pandemic in which individuals were able to withdraw up to A\$20,000 from their government retirement accounts thus increasing liquidity without a change to individual’s lifetime budget constraint. The authors find about one in six eligible people withdrew, and, furthermore, of those who did, most withdrew the maximum amount possible and spent close to half in the first four weeks.

A smaller and negative liquidity shock is studied in [Gelman et al. \(2020\)](#). The authors look at the period during which the U.S. government shut down in 2013 and many federal workers received no (or reduced) pay which was made up to them later. Large negative spending effects—implying an MPC close to 0.5—were observed. However, much of this

effect can be classified as accessing non-standard sources of short-term liquidity, in particular delaying recurring payments such as mortgage and rent payments.

Most natural experiments in this literature, including all those cited thus far, have been related to transitory income shocks. Less is known from the natural experiment literature about the consumption response to permanent shocks to income. [Gelman et al. \(2023\)](#) use changes in gas prices as a proxy for permanent income changes and find, using transaction-level data, that individuals reduce their spending one-for-one with a permanent change in their disposable income. [Gerard and Naritomi \(2021\)](#) look at displaced workers in Brazil who receive a positive transitory shock to income in the form of severance pay, but face a permanent reduction in their lifetime earnings. The authors find that workers increase spending at layoff by 35 percent despite experiencing a 14 percent long-term loss.

While this overview of the natural experiment literature has focused on income *shocks*, there is also a large literature examining the consumption response to *anticipated* income changes. For example, [Ganong and Noel \(2019\)](#) look at spending around the expiration of unemployment benefits, [Hsieh \(2003\)](#) and [Kueng \(2018\)](#) look at spending from the Alaska permanent fund, and [Souleles \(1999\)](#) and [Gelman et al. \(2022\)](#) consider spending around tax refunds. Although standard consumption theory would suggest households' response to anticipated changes should be significantly muted relative to shocks, many of these papers show a large spending change coinciding with the anticipated income change.

4 Elicitation Surveys

The third method that researchers use to understand households' responses to income shocks is to ask them directly in a survey. This has the advantage of being relatively easy to implement without the need to find a natural experiment or get access to data on spending. The method also allows for direct comparison within a household of how their consumption would change in a variety of scenarios, such as in response to permanent versus transitory shocks, or to shocks of different size and sign. However, survey respondents may be unable to know their true behavior.

[Shapiro and Slemrod \(1995, 2003\)](#) started this literature, asking households about their qualitative response to income changes. [Jappelli and Pistaferri \(2014, 2020\)](#) explore heterogeneity, particularly across cash-on-hand holding, using the Italian Survey of Household Income and Wealth which asks respondents for numerical MPCs. [Bunn et al. \(2018\)](#) and [Christelis et al. \(2019\)](#) both explore asymmetry in consumption responses, finding that responses to losses are larger than to gains. [Fuster et al. \(2020\)](#) ask households what they would do with \$500 in the NY Fed survey on consumer expectations. This survey also asked

participants about their responses to news about future income shocks and found that even those with large responses to gains do not respond to news about future gains, although they do cut spending in response to news about future losses. Colarieti et al. (2024) further explore the dynamic spending response to news about future income shocks by asking for spending plans over four quarters. These authors also carry out several cross-validation exercises which suggest that the answers elicited from this type of survey align with actual spending behavior, suggesting these survey methods are of high value.

5 Conclusion

We survey the economics literature of, primarily, the last 20 years, that studies how household consumption responds to income shocks. We group the papers in this literature into three categories: papers that use structural methods, those that exploit natural experiments, and papers that rely on elicitation surveys. The evidence so far suggests that consumption responds more strongly to permanent than to transitory shocks, the sign and size of shocks as well as the horizon over which effects are studied matter, and there is substantial heterogeneity in the responses that cannot be fully explained by observable factors. This is an increasingly busy literature that will remain active in the foreseeable future.

References

- Alan, S., M. Browning, and M. Ejrnæs (2018). Income and consumption: A micro semistructural analysis with pervasive heterogeneity. *Journal of Political Economy* 126(5), 1827–1864.
- Altonji, J. G., F. Hayashi, and L. J. Kotlikoff (1992). Is the extended family altruistically linked? Direct tests using micro data. *American Economic Review* 82(5), 1177–1198.
- Altug, S. and R. A. Miller (1990). Household choices in equilibrium. *Econometrica* 58(3), 543–570.
- Arellano, M., R. Blundell, and S. Bonhomme (2017). Earnings and consumption dynamics: A nonlinear panel data framework. *Econometrica* 85(3), 693–734.
- Arellano, M., R. Blundell, S. Bonhomme, and J. Light (2024). Heterogeneity of consumption responses to income shocks in the presence of nonlinear persistence. *Journal of Econometrics* 240(2).
- Attanasio, O., G. Berloffa, R. Blundell, and I. Preston (2002). From earnings inequality to consumption inequality. *Economic Journal* 112(478), C52–C59.

- Attanasio, O. and S. J. Davis (1996). Relative wage movements and the distribution of consumption. *Journal of Political Economy* 104(6), 1227–1262.
- Attanasio, O., H. Low, and V. Sánchez-Marcos (2005). Female labor supply as insurance against idiosyncratic risk. *Journal of the European Economic Association* 3(2/3), 755–764.
- Attanasio, O. P. and N. Pavoni (2011). Risk sharing in private information models with asset accumulation: Explaining the excess smoothness of consumption. *Econometrica* 79(4), 1027–1068.
- Baker, S., R. A. Farrokhnia, S. Meyer, M. Pagel, and C. Yannelis (2023). Income, liquidity, and the consumption response to the 2020 economic stimulus payments. *Review of Finance* 27(6), 2271–2304.
- Blundell, R., H. Low, and I. Preston (2013). Decomposing changes in income risk using consumption data. *Quantitative Economics* 4(1), 1–37.
- Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *American Economic Review* 98(5), 1887–1921.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review* 106(2), 387–435.
- Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2018). Children, time allocation, and consumption insurance. *Journal of Political Economy* 126(S1), S73–S115.
- Blundell, R. and I. Preston (1998). Consumption inequality and income uncertainty. *Quarterly Journal of Economics* 113(2), 603–640.
- Boehm, J., E. Fize, and X. Jaravel (2022). Five facts about MPCs: Evidence from a randomized experiment. Discussion Paper 18676, Centre for Economic Policy Research.
- Bunn, P., J. Le Roux, K. Reinold, and P. Surico (2018). The consumption response to positive and negative income shocks. *Journal of Monetary Economics* 96, 1–15.
- Busch, C. and A. Ludwig (2021). Higher-order income risk over the business cycle. Unpublished manuscript.
- Chopra, A. (2023). Insurance cyclicalities. Unpublished manuscript.
- Christelis, D., D. Georgarakos, T. Jappelli, L. Pistaferri, and M. van Rooij (2019). Asymmetric consumption effects of transitory income shocks. *Economic Journal* 129(622), 2322–2341.
- Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of Political Economy* 99(5), 957–976.
- Colarieti, R., P. Mei, and S. Stantcheva (2024). The how and why of household reactions to income shocks. NBER Working Paper 32191.

- Commault, J. (2022). Does consumption respond to transitory shocks? Reconciling natural experiments and semistructural methods. *American Economic Journal: Macroeconomics* 14(2), 96–122.
- Crawley, E. (2020). In search of lost time aggregation. *Economics Letters* 189.
- Crawley, E. and A. Kuchler (2023). Consumption heterogeneity: Micro drivers and macro implications. *American Economic Journal: Macroeconomics* 15(1), 314–41.
- De Nardi, M., G. Fella, and G. Paz-Pardo (2019). Nonlinear household earnings dynamics, self-insurance, and welfare. *Journal of the European Economic Association* 18(2), 890–926.
- Deaton, A. and C. Paxson (1994). Intertemporal choice and inequality. *Journal of Political Economy* 102(3), 437–467.
- Fagereng, A., M. B. Holm, and G. J. Natvik (2021). MPC heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics* 13(4), 1–54.
- Falcattoni, E. and V. Nygaard (2021). Acts of Congress and COVID-19: A literature review on the impact of increased unemployment insurance benefits and stimulus checks. FEDS Notes 2021-02-24-2, Board of Governors of the Federal Reserve System.
- Flavin, M. A. (1981). The adjustment of consumption to changing expectations about future income. *Journal of Political Economy* 89(5), 974–1009.
- Fuster, A., G. Kaplan, and B. Zafar (2020). What would you do with \$500? Spending responses to gains, losses, news, loans. *Review of Economic Studies* 88(4), 1760–1795.
- Ganong, P. and P. Noel (2019). Consumer spending during unemployment: Positive and normative implications. *American Economic Review* 109(7), 2383–2424.
- Gelman, M., Y. Gorodnichenko, S. Kariv, D. Koustas, M. D. Shapiro, D. Silverman, and S. Tadelis (2023). The response of consumer spending to changes in gasoline prices. *American Economic Journal: Macroeconomics* 15(2), 129–60.
- Gelman, M., S. Kariv, M. D. Shapiro, and D. Silverman (2022). Rational illiquidity and consumption: Theory and evidence from income tax withholding and refunds. *American Economic Review* 112(9), 2959–91.
- Gelman, M., S. Kariv, M. D. Shapiro, D. Silverman, and S. Tadelis (2020). How individuals respond to a liquidity shock: Evidence from the 2013 government shutdown. *Journal of Public Economics* 189, 103917.
- Gelman, M. and M. Stephens, Jr. (2022). *Lessons learned from Economic Impact Payments during COVID-19*. Brookings.
- Gerard, F. and J. Naritomi (2021). Job displacement insurance and (the lack of) consumption-smoothing. *American Economic Review* 111(3), 899–942.

- Ghosh, A. and A. Theloudis (2023). Consumption partial insurance in the presence of tail income risk. Unpublished manuscript.
- Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review* 87(1), 192–205.
- Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021). What do data on millions of U.S. workers reveal about lifecycle earnings dynamics? *Econometrica* 89(5), 2303–2339.
- Guvenen, F., R. Madera, and S. Ozkan (2023). Consumption dynamics and welfare under non-gaussian earnings risk. Unpublished manuscript.
- Guvenen, F. and A. A. Smith (2014). Inferring labor income risk and partial insurance from economic choices. *Econometrica* 82(6), 2085–2129.
- Hall, R. E. (1978). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy* 86(6), 971–987.
- Hall, R. E. and F. S. Mishkin (1982). The sensitivity of consumption to transitory income: Estimates from panel data on households. *Econometrica* 50(2), 461–481.
- Hamilton, S., G. Liu, J. Miranda-Pinto, and T. Sainsbury (2023). Early pension withdrawal as stimulus. Unpublished manuscript.
- Havranek, T. and A. Sokolova (2020). Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say “probably not”. *Review of Economic Dynamics* 35, 97–122.
- Hayashi, F., J. Altonji, and L. Kotlikoff (1996). Risk-sharing between and within families. *Econometrica* 64(2), 261–294.
- Heathcote, J., K. Storesletten, and G. L. Violante (2008). Insurance and opportunities: A welfare analysis of labor market risk. *Journal of Monetary Economics* 55(3), 501–525.
- Heathcote, J., K. Storesletten, and G. L. Violante (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review* 104(7), 2075–2126.
- Hryshko, D. and I. Manovskii (2022). How much consumption insurance in the U.S.? *Journal of Monetary Economics* 130, 17–33.
- Hsieh, C.-T. (2003). Do consumers react to anticipated income changes? Evidence from the Alaska permanent fund. *American Economic Review* 93(1), 397–405.
- Hyslop, D. R. (2001). Rising U.S. earnings inequality and family labor supply: The covariance structure of intrafamily earnings. *American Economic Review* 91(4), 755–777.
- Jappelli, T. and L. Pistaferri (2010). The consumption response to income changes. *Annual Review of Economics* 2(1), 479–506.

- Jappelli, T. and L. Pistaferri (2014). Fiscal policy and MPC heterogeneity. *American Economic Journal: Macroeconomics* 6(4), 107–36.
- Jappelli, T. and L. Pistaferri (2020). Reported MPC and unobserved heterogeneity. *American Economic Journal: Economic Policy* 12(4), 275–97.
- Jessen, R. and J. König (2023). Hours risk and wage risk: repercussions over the life cycle. *The Scandinavian Journal of Economics* 125(4), 956–996.
- Johnson, D. S., J. A. Parker, and N. S. Souleles (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review* 96(5), 1589–1610.
- Kaplan, G. and G. L. Violante (2010). How much consumption insurance beyond self-insurance? *American Economic Journal: Macroeconomics* 2(4), 53–87.
- Karger, E. and A. Rajan (2020). Heterogeneity in the marginal propensity to consume: Evidence from Covid-19 stimulus payments. Working Paper Series WP-2020-15, Federal Reserve Bank of Chicago.
- Krueger, D. and F. Perri (2006). Does income inequality lead to consumption inequality? Evidence and theory. *Review of Economic Studies* 73(1), 163–193.
- Kueng, L. (2018). Excess sensitivity of high-income consumers. *Quarterly Journal of Economics* 133(4), 1693–1751.
- Lewis, D. J., D. Melcangi, and L. Pilossoph (2019). Latent heterogeneity in the marginal propensity to consume. Staff Report 902, Federal Reserve Bank of New York.
- Low, H., C. Meghir, and L. Pistaferri (2010). Wage risk and employment risk over the life cycle. *American Economic Review* 100(4), 1432–67.
- Mace, B. J. (1991). Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy* 99(5), 928–956.
- Madera, R. (2019). The consumption response to tail earnings shocks. Unpublished manuscript.
- Meghir, C. and L. Pistaferri (2011). Earnings, consumption and life cycle choices. Volume 4 of *Handbook of Labor Economics*, pp. 773–854. Elsevier.
- Misra, K., V. Singh, and Q. Zhang (2022). Impact of stay-at-home-orders and cost-of-living on stimulus response: Evidence from the CARES Act. *Marketing Science* 41(2), 211–229.
- Misra, K. and P. Surico (2014). Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics* 6(4), 84–106.
- Orchard, J., V. A. Ramey, and J. F. Wieland (2023). Micro MPCs and macro counterfactuals: The case of the 2008 rebates. NBER Working Paper 31584.

- Parker, J. A., J. Schild, L. Erhard, and D. Johnson (2022). Economic impact payments and household spending during the pandemic. NBER Working Paper 30596.
- Parker, J. A., N. S. Souleles, D. S. Johnson, and R. McClelland (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review* 103(6), 2530–53.
- Primiceri, G. and T. van Rens (2009). Heterogeneous life-cycle profiles, income risk and consumption inequality. *Journal of Monetary Economics* 56(1), 20–39.
- Sabelhaus, J., D. Johnson, S. Ash, D. Swanson, T. I. Garner, J. Greenlees, and S. Henderson (2014). *Is the Consumer Expenditure Survey Representative by Income?*, pp. 241–262. University of Chicago Press.
- Sahm, C. R., M. D. Shapiro, and J. Slemrod (2010). Household response to the 2008 tax rebate: Survey evidence and aggregate implications. *Tax Policy and the Economy* 24(1), 69–110.
- Sahm, C. R., M. D. Shapiro, and J. Slemrod (2012). Check in the mail or more in the paycheck: Does the effectiveness of fiscal stimulus depend on how it is delivered? *American Economic Journal: Economic Policy* 4(3), 216–50.
- Sargent, T. J. (1978). Rational expectations, econometric exogeneity, and consumption. *Journal of Political Economy* 86(4), 673–700.
- Shapiro, M. D. and J. Slemrod (1995). Consumer response to the timing of income: Evidence from a change in tax withholding. *American Economic Review* 85(1), 274–283.
- Shapiro, M. D. and J. Slemrod (2003). Consumer response to tax rebates. *American Economic Review* 93(1), 381–396.
- Shapiro, M. D. and J. Slemrod (2009). Did the 2008 tax rebates stimulate spending? *American Economic Review* 99(2), 374–79.
- Souleles, N. S. (1999). The response of household consumption to income tax refunds. *American Economic Review* 89(4), 947–958.
- Storesletten, K., C. I. Telmer, and A. Yaron (2004). Consumption and risk sharing over the life cycle. *Journal of Monetary Economics* 51(3), 609–633.
- Theloudis, A. (2017). *Essays in Consumption Inequality and the Allocation of Household Resources*. Ph. D. thesis, University College London.
- Theloudis, A. (2021). Consumption inequality across heterogeneous families. *European Economic Review* 136.
- Townsend, R. M. (1994). Risk and insurance in village India. *Econometrica* 62(3), 539–591.
- Wu, C. and D. Krueger (2021). Consumption insurance against wage risk: Family labor supply and optimal progressive income taxation. *American Economic Journal: Macroeconomics* 13(1), 79–113.